

Emerging Optimal Control Models and Solvers for Interconnected Natural Gas and Electricity Networks

Nai-Yuan Chiang and Victor M. Zavala

Abstract This chapter reviews emerging optimal control models for interconnected natural gas and electricity networks and discusses economic drivers motivating the development of such models. We also review computational patterns and structures arising in these models and assess the potential and limitations of state-of-the-art optimization solvers.

1 Motivation

Natural gas and power grid infrastructures are becoming increasingly interdependent. A major factor driving this situation is the increasing deployment of gas-fired power plants [1]. These plants are modular and less capital intensive compared with large, centralized generation facilities running on nuclear and coal fuel sources. In addition, gas-fired plants are more flexible and can quickly ramp up and down their power output. This flexibility becomes an asset as the share of intermittent solar and wind power increases. Moreover, the high availability of gas resulting from new fracking technologies has led to lower prices, making gas-fired plants economically more attractive.

An important feature of gas-fired generation (compared with other generation technologies) is that large amounts of fuel must be transported to power generation facilities in gaseous form through a sophisticated network that spans thousands

Nai-Yuan Chiang
Mathematics and Computer Science Division
Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439, USA
e-mail: nychiang@mcs.anl.gov

Victor M. Zavala
Department of Chemical and Biological Engineering
University of Wisconsin-Madison, 1415 Engineering Drive, Madison, WI 53706, USA
e-mail: victor.zavala@wisc.edu

of miles. A key advantage of having such a network infrastructure is that significant amounts of gas can be stored inside the pipelines. The stored gas is distributed spatially in the network and is normally referred to as *line-pack* [8]. Line-pack is used by pipeline operators to modulate variations of gas demands at multiple spatial points in intraday operations. Some of the strongest variations in gas demands are the result of on-demand start-up and shut-down of gas-fired power plants [22]. Modulating these variations is challenging because the fast release of line-pack at multiple simultaneous locations can trigger complex spatiotemporal responses that propagate hundreds to thousands of miles and that can take hours to stabilize. Therefore, line-pack management is performed by using sophisticated optimal control and pipeline simulation tools. These automation tools orchestrate the operation of a multitude of compressor stations distributed throughout the system with the objectives of satisfying demands, maintaining pressure levels, and minimizing compression costs [23].

An important issue faced by gas-fired power plants is that they compete for natural gas with industrial facilities and with local distribution companies (LDCs) that supply gas to urban areas. Therefore, natural gas cannot be guaranteed to be available at each power generation facility at all times. This limitation is particularly evident during the winter season when residential and office buildings require large amounts of gas for heating. An extreme manifestation of this issue was observed during the polar vortex of 2014 in which sustained low temperatures in the Midwest region of the U.S. led to high gas demands in urban areas and to equipment failures^{1,2}. These factors resulted in widespread shortages of natural gas in places as remote as California, Massachusetts, and Texas. These gas shortages in turn resulted in lost electrical generation capacity totaling 35 GW. At a value of lost load of 5,000 \$/MWh, shortages of this magnitude represent economic losses of 175 million \$/hr. The New England area alone lost 1.5 GW of power generation capacity [2]. The polar vortex also exposed market inefficiencies resulting from the increasing interaction between grid and gas systems. In particular, gas-fired plants required significant uplift payments from the independent system operators (ISOs). These payments compensated the power plants for the lost revenue resulting from the inability of the gas infrastructure to deliver fuel. These operational and economic issues question the ability of the gas network infrastructure to sustain additional gas-fired generation. Consequently, it is important to investigate the economic and flexibility of natural gas infrastructures by developing detailed physical models and advanced optimal control tools.

Optimal control formulations of natural gas networks using high-resolution dynamic models have been reported by several researchers [26, 8, 27]. These studies, however, do not consider coordination with power grid networks and treat power plants as exogenous uncertain disturbances. By coupling gas and electric models one can better understand the enhanced flexibility and resiliency that can be gained by coordination. Researchers have also developed models of interconnected power

¹ <http://tinyurl.com/pvouhym>

² <http://tinyurl.com/no5v2za>

grid and natural gas networks but the dynamics of natural gas systems are often neglected [14, 5, 4]. Steady-state models cannot capture line-pack storage dynamics and thus significantly misrepresent the flexibility of the system in real-time operations. Consequently, steady-state models are more appropriate for long-term planning studies [25]. In this respect, our model seeks to better capture the flexibility provided by line-pack in real-time operations. Recent studies have also reported models and strategies for co-optimization of gas and power grid transmission systems using detailed dynamic gas models. The studies in [16, 17] use full-resolution models but focus on small synthetic models to assess economic improvements resulting from coordination and to evaluate the impact of using dynamic over steady-state gas pipeline models. The studies in [10, 21] focus on the Great Britain network and provide more in-depth analyses. In particular, the study in [10] presents a multi-time period model to study the effects of gas terminal (supply) failures on the integrated gas-electric system. The model, however, uses simplifications to address computational complexity; in particular, an aggregated Great Britain model with 16 buses is used and the gas system dynamics are only captured at the daily time scale (dynamics in intra-day operations are ignored). The model proposed in [21] studies the effect of wind power adoption levels on gas generation and demonstrates that line-pack can limit system performance during periods with low wind generation. The simplified 16-bus Great Britain network is also used in this study, the gas network is simplified by aggregating parallel branches, and gas dynamics are ignored. None of the studies reported in the literature discuss computational efficiency and scalability issues. In this chapter we review modern optimal control formulations of natural gas networks and discuss the economic drivers motivating such applications. We also provide a review of existing computational capabilities that enable the solution of these complex optimization problems and provide pointers to open challenges. Our objective is to focus on real-sized networks and to capture full physical and spatiotemporal resolutions. This can enable us to discover non-intuitive physical behavior that can inform ISOs and gas pipeline operators. Focusing on real-sized systems will also enable us to assess the limits of state-of-the-art optimization algorithms and to identify scalability bottlenecks.

The chapter is structured as follows: Section 2 presents elements of a typical optimal control model for natural gas networks and discusses interconnection with power grid models and derivation of stochastic formulations. Section 3 discusses economic and resiliency insights that can be gained with detailed optimal control models and Section 4 discusses computational solution strategies and open challenges. Concluding remarks are provided in Section 5.

2 Optimal Control Formulations

In this section we present typical components of a optimal control formulation for gas networks and provide pointers to suitable extensions.

2.1 Transport Equations

We use the following isothermal form of the continuity and momentum equations:

$$\frac{\partial p_\ell(x, \tau)}{\partial \tau} + \frac{ZRT}{A_\ell} \frac{\partial f_\ell(x, \tau)}{\partial x} = 0 \quad (1a)$$

$$\frac{1}{A_\ell} \frac{\partial f_\ell(x, \tau)}{\partial \tau} + \frac{\partial p_\ell(x, \tau)}{\partial x} + \frac{8\lambda_\ell}{\pi^2 D_\ell^5} \frac{f_\ell(x, \tau) |f_\ell(x, \tau)|}{\rho_\ell(x, \tau)} = 0 \quad (1b)$$

Here, $\tau \in [0, N] := \mathcal{T}$ is the time dimension with final time N and $x \in \mathcal{X}_\ell := [0, L_\ell]$ is the axial dimension with length L_ℓ . The link diameters are denoted as D_ℓ and the friction coefficients as λ_ℓ . The states of the link are the gas density $\rho_\ell(x, \tau)$, speed $v_\ell(x, \tau)$, and pressure $p_\ell(x, \tau)$. The gas pressure and density are related as

$$\frac{p_\ell(x, \tau)}{\rho_\ell(x, \tau)} = ZRT, \quad (2)$$

where R is the universal gas constant and T is temperature. For a detailed non-isothermal formulation the reader can refer to [3]. The pipeline links are connected through a network comprising a set \mathcal{N} of nodes, a set \mathcal{S} of supply flows, and a set \mathcal{D} of demand flows. For each node $n \in \mathcal{N}$ we define the set of inlet and outlet links, $\mathcal{L}_n^{snd} := \{\ell \mid snd(\ell) = n\}$, $\mathcal{L}_n^{rec} := \{\ell \mid rec(\ell) = n\}$. Here, $snd(\ell) \in \mathcal{N}$ is the start node of link ℓ and $rec(\ell) \in \mathcal{N}$ is the end node. We assume that the direction of the flow is given. For formulations that allow for flow reversals the reader is referred to [6]. We define $dn(j) \in \mathcal{N}$ as the node at which the demand flow $d_j(\tau)$ is located and $sn(i) \in \mathcal{N}$ as the node at which the supply flow $s_i(\tau)$ is located. Accordingly, we define the sets $\mathcal{S}_n := \{j \in \mathcal{S} \mid sn(j) = n\}$ and $\mathcal{D}_n := \{j \in \mathcal{D} \mid dn(j) = n\}$ for each node $n \in \mathcal{N}$. For each node $n \in \mathcal{N}$ we also define pressures $\theta_n(\cdot)$.

For modeling convenience we lift the network system by introducing *dummy* inlet flows for each link $f_\ell^{in}(\cdot)$, respectively; and outlets $f_\ell^{out}(\cdot)$. Using these definitions we can express mass balances at the nodes as

$$\sum_{\ell \in \mathcal{L}_n^{rec}} f_\ell^{out}(\tau) - \sum_{\ell \in \mathcal{L}_n^{snd}} f_\ell^{in}(\tau) + \sum_{i \in \mathcal{S}_n} s_i(\tau) - \sum_{j \in \mathcal{D}_n} d_j(\tau) = 0, \quad n \in \mathcal{N} \quad (3)$$

We split the set of links \mathcal{L} into subsets of passive \mathcal{L}_p links and active links \mathcal{L}_a . For the active links we define the *boost pressures* $\Delta\theta_\ell(\cdot)$ which are the additional (non-negative) pressures introduced by the compressor located at the inlet (sending) node of the link. For the passive links, there is no compression. The boundary conditions for the link pressures are given by

$$p_\ell(L_\ell, \tau) = \theta_{rec(\ell)}(\tau), \quad \ell \in \mathcal{L} \quad (4a)$$

$$p_\ell(0, \tau) = \theta_\ell^{dis}(\tau), \quad \ell \in \mathcal{L}. \quad (4b)$$

The boundary conditions for the link flows are,

$$f_\ell(0, \tau) = f_\ell^{in}(\tau), \ell \in \mathcal{L} \quad (5a)$$

$$f_\ell(L_\ell, \tau) = f_\ell^{out}(\tau), \ell \in \mathcal{L}. \quad (5b)$$

The discharge pressures of the compressors (located at the inlet of the active links) are given by,

$$\theta_\ell^{dis}(\tau) = \theta_{snd(\ell)}(\tau) + \Delta\theta_\ell(\tau), \ell \in \mathcal{L}_a. \quad (6a)$$

For the passive links we simply have that,

$$\theta_\ell^{dis}(\tau) = \theta_{snd(\ell)}(\tau), \ell \in \mathcal{L}_p. \quad (7)$$

The total compression power consumed in the active links is given by

$$P_\ell(\tau) = c_p f_\ell^{in}(\tau) T \left(\left(\frac{\theta_\ell^{dis}(\tau)}{\theta_{rec(\ell)}(\tau)} \right)^\beta - 1 \right), \ell \in \mathcal{L}_a. \quad (8)$$

As can be seen, the transport equations are highly nonlinear. Additional complexity can be added to the model by consider control valves and more complex compressor configurations [26, 18].

2.2 Constraints

We consider the following constraints on available compressor power, suction and discharge pressures, and demand delivery pressures,

$$P_\ell^L \leq P_\ell(\tau) \leq P_\ell^U, \ell \in \mathcal{L}_a \quad (9a)$$

$$\theta_\ell^{suc,L} \leq \theta_{snd(\ell)}(\tau) \leq \theta_\ell^{suc,U}, \ell \in \mathcal{L}_a \quad (9b)$$

$$\theta_\ell^{dis,L} \leq \theta_\ell^{dis}(\tau) \leq \theta_\ell^{dis,U}, \ell \in \mathcal{L}_a \quad (9c)$$

$$\theta_j^{dem,L} \leq \theta_{dn(j)}(\tau) \leq \theta_j^{dem,U}, j \in \mathcal{D}. \quad (9d)$$

We also have the implicit physical bounds $f_\ell^{in}(\cdot), f_\ell^{out}(\cdot), P_\ell(\cdot) \geq 0$ and we assume that the pressures (or inlet flows but not both) at the supply points are given:

$$\theta_{sn(j)}(\tau) = \text{given} \quad \text{or} \quad s_j(\tau) = \text{given}, \quad j \in \mathcal{S}. \quad (10)$$

The gas demands are bounded by the given targets $d_j^{target}(\cdot)$,

$$0 \leq d_j(\tau) \leq d_j^{target}(\tau), j \in \mathcal{D}. \quad (11)$$

We enforce periodicity of the line-pack in each pipeline,

$$\int_0^{L_\ell} f_\ell(x, T) dx \geq \int_0^{L_\ell} f_\ell(x, 0) dx, \quad \ell \in \mathcal{L}_a. \quad (12)$$

Without this constraint, the system will tend to deplete line-pack in order to minimize compressor power and this can make operation in the following day infeasible. The periodicity constraint thus provides a mechanism to deal with the finite horizon of the optimal control problem (OCP) and ensure recursive feasibility [24].

2.3 Initial State

We assume that the system is at steady-state at the initial time $\tau = 0$. Such steady-state satisfies the steady-state transport equations

$$\frac{ZRT}{A_\ell} \frac{\partial f_\ell(x, \tau)}{\partial x} = 0 \quad (13a)$$

$$\frac{\partial p_\ell(x, \tau)}{\partial x} + \frac{8\lambda_\ell}{\pi^2 D_\ell^5} \frac{f_\ell(x, \tau) |f_\ell(x, \tau)|}{\rho_\ell(x, \tau)} = 0 \quad (13b)$$

One can use the boundary conditions defined at $\tau = 0$ and the assumption of given supply pressures (or flows) at $\tau = 0$ to prove that fixing boost pressures $\Delta \theta_\ell(0)$ and demand flows $d_j(0)$ fully defines the initial states [27].

2.4 Objective Function

The cost function can be set as a combination of gas supply, compression cost, and gas demand delivery,

$$\phi_{gas} := \int_0^N \left(\sum_{i \in \mathcal{S}} \alpha_i^s s_i(\tau) + \sum_{\ell \in \mathcal{L}_a} \alpha_\ell^P P_\ell(\tau) - \sum_{j \in \mathcal{D}} \alpha_j^d d_j(\tau) \right) d\tau \quad (14)$$

The first term is the supply cost, the second term in the cost function is the compressor power with cost α_ℓ^P , and the third term is the total value of served demand. The demand term seeks to maximize the served demand (this explains the negative sign) and the parameter α_j^d can be interpreted as the value of the served demand. Consequently, the cost function is an analog of the social welfare used in electricity market clearing formulations [20]. From a control stand-point, the demand term can also be interpreted as penalty for unserved demand. In particular, when the value α_j^d is high relative to compression power and supply, the control system will tend to push the demands $d_j(\cdot)$ to the targets $d_j^{target}(\cdot)$ defined in the constraints (11). The parameter α_j^d can also be used to prioritize certain demand locations (e.g., residential over power plants) and can be used to model gas bids provided by gas-fired power

plants. We highlight that the cost function (14) is linear and this tends to introduce ill-conditioning in the Hessian matrix and slow down convergence. This is an issue arising in optimal control models with economic objectives and is exacerbated in distributed systems [28]. To ameliorate this issue, it is also possible to define a cost function of the form

$$\varphi_{gas} := \int_0^N \left(\sum_{i \in \mathcal{S}} \alpha_i^s s_i(\tau) + \sum_{\ell \in \mathcal{L}_a} \alpha_\ell^P P_\ell(\tau) + \sum_{j \in \mathcal{D}} \alpha_j^d (d_j(\tau) - d_j^{target}(\tau))^2 \right) d\tau. \quad (15)$$

Here, the quadratic term in the objective is a tracking term that seeks to bring the demands to the targets. An advantage of this formulation is that it adds positive curvature to the cost function and this can aid computational performance. A disadvantage of this formulation is that the parameter α_j^d does not necessarily have an economic interpretation and, consequently, it might be difficult to tune.

The optimal control model for the gas side can be summarized as:

$$\begin{aligned} \min \quad & \varphi_{gas} \\ \text{s.t.} \quad & (1) - (13). \end{aligned}$$

One can use a structural model analysis to prove that the number of *degrees of freedom* of the problem corresponds to that of the control trajectories $\Delta \theta_\ell(\tau)$, $\ell \in \mathcal{L}_a$, $\tau \in \mathcal{T}$ and $d_j(\tau)$, $j \in \mathcal{D}$, $\tau \in \mathcal{T}$ [27].

2.5 Integrated with Gas-Electric Formulations

A major source of uncertainty in natural gas systems are the large gas demands resulting from power plants [22]. At the same time, gas-fired power plants cannot guarantee that the entire amount of gas requested can be delivered at all times and this uncertainty affects ISO operations. Consequently, it is natural to believe that better communication/coordination between natural gas and power grid infrastructures can significantly help mitigate uncertainty on both sides. Moreover, we can expect that by combining control flexibility through coordinate we can increase overall resiliency. The question is, however, by how much? To answer this question we can develop integrated gas-electric optimal control models.

Economic dispatch is an OCP that is solved by ISOs to balance supply and demand and to price electricity in intraday operations. We formulate the dispatch problem as the continuous-time OCP shown in (16). For a more detailed explanation of the model the reader is referred to [13]. Here, we use a notation similar to that used in the description of the gas system to highlight similarities and differences between gas and electric models. We highlight differences in notation when appropriate. We define the sets of electricity suppliers (i.e., power plants or generators) as \mathcal{S} , the set of electrical loads as \mathcal{D} , the set of network nodes as \mathcal{N} , and the set of transmission

lines as \mathcal{L} . For each link $\ell \in \mathcal{L}$ we denote $snd(\ell) \in \mathcal{N}$ as the sending node and $rec(\ell) \in \mathcal{N}$ as the receiving node. We define $\mathcal{S}_g \subseteq \mathcal{S}$ as the subset of power plants that are fired by natural gas. The rest of the suppliers are either thermal or renewable power suppliers. We define $\mathcal{S}_n \subseteq \mathcal{S}$ as the subset of suppliers connected to node n and $\mathcal{D}_n \subseteq \mathcal{D}$ as the subset of loads connected to node n . We define $\mathcal{L}_n^{snd} \subseteq \mathcal{L}$ as the set of links originating from node n and $\mathcal{L}_n^{rec} \subseteq \mathcal{L}$ as the set of lines ending in node n .

$$\min \quad \varphi_{grid} := \int_0^N \left(\sum_{i \in \mathcal{S}} \alpha_i^s s_i(\tau) - \sum_{j \in \mathcal{D}} \alpha_j^d d_j^{grid}(\tau) \right) d\tau \quad (16a)$$

$$\text{s.t.} \quad \frac{ds_i(\tau)}{d\tau} = r_i(\tau), \quad i \in \mathcal{S} \quad (16b)$$

$$\sum_{\ell \in \mathcal{L}_n^{rec}} f_\ell(\tau) - \sum_{\ell \in \mathcal{L}_n^{snd}} f_\ell(\tau) + \sum_{i \in \mathcal{S}_n} s_i(\tau) - \sum_{j \in \mathcal{D}_n^{grid}} d_j^{grid}(\tau) = 0, \quad n \in \mathcal{N} \quad (16c)$$

$$f_\ell(\tau) = \beta_\ell (\theta_{snd(\ell)}(\tau) - \theta_{rec(\ell)}(\tau)), \quad \ell \in \mathcal{L} \quad (16d)$$

$$\underline{f}_\ell \leq f_\ell(\tau) \leq \bar{f}_\ell, \quad \ell \in \mathcal{L} \quad (16e)$$

$$\underline{\theta}_n \leq \theta_n(\tau) \leq \bar{\theta}_n, \quad n \in \mathcal{N} \quad (16f)$$

$$\underline{s}_i \leq s_i(\tau) \leq \bar{s}_i, \quad i \in \mathcal{S} \quad (16g)$$

$$\underline{r}_i \leq r_i(\tau) \leq \bar{r}_i, \quad i \in \mathcal{S} \quad (16h)$$

$$0 \leq d_j^{grid}(\tau) \leq d_j^{grid, target}(\tau), \quad i \in \mathcal{S} \quad (16i)$$

$$d_i^{gas, grid}(\tau) = \eta_i \cdot s_i(\tau), \quad i \in \mathcal{S}_g. \quad (16j)$$

The time profiles for power generation, delivered loads, flows, and voltage angles are denoted as $s_i(\cdot)$, $i \in \mathcal{S}$; $d_j^{grid}(\cdot)$, $j \in \mathcal{D}^{grid}$; $f_\ell(\cdot)$, $\ell \in \mathcal{L}$; and θ_n , $n \in \mathcal{N}$, respectively. The time profiles for target loads are denoted as $d_j^{grid, target}(\cdot)$. The power flows are bounded by the capacity limits $\underline{f}_\ell, \bar{f}_\ell$, the angles are bounded by $\underline{\theta}_n, \bar{\theta}_n$, and the supply flows are bounded by $\underline{s}_i, \bar{s}_i$. The supply change rates are given by $r_i(\cdot)$, $i \in \mathcal{S}$ and are bounded by the ramp limits $\underline{r}_i, \bar{r}_i$.

The cost function (16a) is the negative social welfare [20]. The generation and demands costs are α_i^s , $i \in \mathcal{S}$ and α_j^d , $j \in \mathcal{D}^{grid}$, respectively. In the case of inelastic demands the demand costs are typically set to the value of lost load (VOLL). This value is typically in the range 1,000-10,000 \$/MWh [9]. This parameter has similar interpretation to the one used in the gas side.

The dual variables of the network balance equation (16c) are the locational marginal prices, which we denote as $\pi_n(\cdot)$, $n \in \mathcal{N}$. Equations (16d) are the DC power flow relationships, equations (16f)-(16h) are bounds for voltage angles, generation, and ramps; respectively.

We define $d_i^{gas, grid}(\cdot)$, $i \in \mathcal{S}_g$ as the gas demands originating from the gas-fired plants and η_i , $i \in \mathcal{S}_g$ as the heat rates of the different power plants. The *heat rate* is a measure of the conversion efficiency of a power plant installation and is defined

as the amount of fuel (in BTUs) needed to produce a KWh of electrical energy. The ideal heat rate is 3,412 BTU/KWh, which indicates a one-to-one conversion between fuel and electrical energy (efficiency of 100%). Consequently, the smaller the heat rate, the more efficient the technology. Different gas-fired plants can have different heat rates (combined cycle plants have much smaller heat rates than do simple-cycle plants). In fact, the energy information administration (EIA) reports that the average heat rate for gas-fired plants has been reducing steadily from 9,207 BTU/KWh in 2003 (efficiency of 37%) to 7,948 BTU/KWh in 2013 (efficiency of 42.9%) because of the adoption of new technologies. In comparison, the average heat rate for coal power plants was reported to be 10,459 BTU/KWh (efficiency of 32%) in 2013 and this has remained at similar levels since 2003³.

The coupling between gas and power grid infrastructures is given by the following constraints:

$$d_j^{target}(\tau) = d_{gn(j)}^{gas,grid}(\tau) + d_j^{base}(\tau), \quad j \in \mathcal{D} \quad (17a)$$

$$d_{gn(j)}^{gas,grid}(\tau) \leq d_j(\tau) - d_j^{base}(\tau), \quad j \in \mathcal{D}. \quad (17b)$$

The first constraint states that the gas demand targets for the gas infrastructure are given by the gas demands of the gas-fired power plants plus an exogenous base gas demand that arises from industrial facilities and/or LDCs serving urban areas. Here, $gn(j) \in \mathcal{S}_g$ denotes the power plant corresponding to the gas demand $j \in \mathcal{D}$. The second constraint states that the gas used by the power plants cannot physically exceed the gas delivered by the gas infrastructure. We use this second constraint to model situations in which the gas infrastructure is physically constrained and thus the delivered demand $d_j^{gas}(\cdot)$ cannot match the target demand $d_{gn(j)}^{gas,target}(\cdot)$. The simultaneous solution of the grid and gas models together with the coupling constraints (17) gives the *coordinated* dispatch model.

To compare the performance of a coordinated gas-electric system, we consider an *uncoordinated* setting in which the two infrastructures do not dispatch jointly (as is currently done). This is simulated by first solving the economic dispatch problem for the power grid (16) to set the *predicted* generation $s_i(\cdot)$, $i \in \mathcal{S}$, the natural gas demand targets $d_{gn(j)}^{gas,grid}(\cdot)$, $j \in \mathcal{D}_g$, and the predicted locational marginal prices $\pi_n(\cdot)$, $n \in \mathcal{N}$. Having the gas demand targets, we solve the gas dispatch problem to maximize the gas delivered and minimize compression costs. The solution of this problem sets the *realized* gas demands delivered to the gas-fired power plants $d_j^{gas}(\cdot) - d_j^{gas,base}(\cdot)$. Because the realized gas demands might not be able to match the power grid targets, we solve the economic dispatch problem for the power grid (16) again to determine the realized generation schedule and locational marginal prices corresponding given the realized delivered gas demands. We denote the realized power generation schedules as $s_i^{real}(\cdot)$ and the prices as $\pi_n^{real}(\cdot)$. Differences between the target and delivered gas demands will introduce a difference between the predicted and realized generation schedules and prices. When the gas-fired power

³ http://www.eia.gov/electricity/annual/html/epa_08_01.html

plants cannot obtain the total gas requested, they will need to curtail power and must pay for the unserved electricity generation at the realized price [19]. In such a case, the revenue for the power plants is given by

$$\mathcal{R}_i := \int_0^N \left(\pi_{sn(i)} s_i(\tau) + s_i^{real}(\tau) (\pi_{sn(i)}^{real}(\tau) - \pi_{sn(i)}(\tau)) - \alpha_i^s s_i^{real}(\tau) \right) d\tau, \quad i \in \mathcal{S}_g. \quad (18)$$

Here, $sn(i) \in \mathcal{N}$ denotes the node at which supplier $i \in \mathcal{S}_g$ is connected to. The total revenue for the gas-fired generators is denoted by $\mathcal{R} = \sum_{i \in \mathcal{S}_g} \mathcal{R}_i$.

When the requested and delivered gas demands coincide we have that the predicted and realized generation schedules coincide (i.e., the predicted generation schedule is feasible to the gas system). Consequently, we have that $\pi_{sn(i)}^{real}(\cdot) = \pi_{sn(i)}(\cdot)$ and the revenue reduces to

$$\mathcal{R}_i = \int_0^N \left(\pi_{sn(i)}(\tau) s_i(\tau) - \alpha_i^s s_i^{real}(\tau) \right) d\tau, \quad i \in \mathcal{S}_g. \quad (19)$$

After space-time discretization (for a review on discretization techniques the reader is referred to [3, 7]), we can represent the coupled gas-electric problem as the nonlinear program (NLP):

$$\min \varphi_{grid}(w_{grid}) + \varphi_{gas}(w_{gas}) \quad (20a)$$

$$\text{s.t. } c_{grid}(w_{grid}) \geq 0, \quad (\lambda_{grid}) \quad (20b)$$

$$c_{gas}(w_{gas}) \geq 0, \quad (\lambda_{gas}) \quad (20c)$$

$$\Pi_{gas} w_{gas} + \Pi_{grid} w_{grid} = 0, \quad (\lambda). \quad (20d)$$

Here, w_{grid} are all the variables in the grid side and w_{gas} are all the variables in the gas side. This problem decomposes if the coupling constraints (20d) are removed. The coupling constraints correspond to (17) (after introducing slack variables) and we note that all that is needed to form the coupling constraints are the matrices Π_{gas} and Π_{grid} . These are trivial matrices (containing only zeros and ones).

2.6 Stochastic Formulations

In the presence of uncertainty, both gas and grid sides will seek to make decisions in anticipation of the future. We can formulate this decision process as a two-stage problem of the form in (21a). Here, $\omega \in \Omega$ is the scenario realization, $\mathbb{E}[\cdot]$ denotes the expectation, and constraints (21e) model nonanticipativity conditions for the grid and gas sides. The expected value form can also be used to model conditional-value-at-risk functions [27]. The nonanticipativity conditions separate the variables into here-and-now decisions (decisions made before uncertainty is revealed) and recourse variables (decisions made after uncertainty is revealed). Here-and-now de-

cisions in the gas side can be the allocation of line-pack while in the grid side are forward generation schedules. Recourse decisions on the gas side can be corrections of compressor power to mitigate variability of gas demand while on the power grid side can be corrections on generation to mitigate wind power fluctuations. For more details on stochastic formulations for gas and power grid networks the reader is referred to [27, 20].

$$\min \mathbb{E}[\varphi_{grid}(w_{grid,\omega}) + \varphi_{gas}(w_{gas,\omega})] \quad (21a)$$

$$\text{s.t. } c_{grid,\omega}(w_{grid,\omega}) \geq 0, \quad \omega \in \Omega, \quad (\lambda_{grid,\omega}) \quad (21b)$$

$$c_{gas,\omega}(w_{gas,\omega}) \geq 0, \quad \omega \in \Omega, \quad (\lambda_{gas,\omega}) \quad (21c)$$

$$\Pi_{gas,\omega} w_{gas,\omega} + \Pi_{grid,\omega} w_{grid,\omega} = 0, \quad \omega \in \Omega \quad (\lambda_\omega) \quad (21d)$$

$$\sum_{\omega \in \Omega} \Pi_\omega(w_{gas,\omega}, w_{grid,\omega}) = 0. \quad (\lambda) \quad (21e)$$

The solution the coupled stochastic problem (21a) gives rise to here-and-now control policies. Uncertainty can model equipment failures or externalities such as wind and solar power. When the systems are decoupled, uncertainty can also model the interface variables between infrastructures (i.e., gas demands from power plants $d_j^{target}(\cdot)$). We also recall that a deterministic formulation computes the policies using average information of the uncertain data and we recall that a wait-and-see (perfect information) policy can be obtained by dropping the nonanticipativity constraints.

3 Economic and Resiliency Issues

In this section we highlight economic and resiliency gains that can be achieved by coordination. We first compare how a stochastic formulation that uses uncertainty information from the power grid can significantly aid resiliency of pipeline operations. To demonstrate this, we use a long pipeline system with 13 nodes and 12 pipelines and 10 compressors that spans 1,600 km [27].

In Figure 1 we present the axial flow profiles for a pipeline system for three demand scenarios. The gray lines are the axial profiles during the charging phase of line-pack in the system. The axial profiles in this charging phase are the same for all scenarios because these are non-anticipative decisions. Note that demand is satisfied at each time step for all scenarios. From Figure 1 we also see that the optimal here-and-now (HN) policy consists on progressively accumulate line-pack toward the end of the system reflected by a large increase of flow close to the demand node. Once uncertainty is revealed, the system takes three different paths. The total compression energy for the low, medium, and high demand scenarios is 92742 kWh, 99783 kWh, and 113672 kWh, respectively.




Fig. 1 Optimal flow profiles for low (left), medium (middle), and high (right) demand scenarios. Stochastic (here-and-now) policy.




Fig. 2 Optimal axial flow profiles for low (left), medium (middle), and high (right) demand scenarios. Deterministic policy.




Fig. 3 Optimal axial flow profiles for low (left), medium (middle), and high (right) demand scenarios. Wait-and-see (perfect information) policy.

We now compare the performance of the deterministic and wait-and-see (WS) policies under perfect information. The optimal flow profiles for the deterministic formulation are presented in Figure 2. As can be seen, the policy does not build as much inventory as does the HN solution presented in Figure 1. As a result, while the low and medium demand scenarios are feasible, the system struggles to satisfy the high demand scenario. In fact, we have found that the demand needs to be curtailed for the system to remain feasible. This high-stress behavior is also reflected in highly volatile flow profiles resulting from aggressive recourse actions.

The axial flow profiles for the WS solution are presented in Figure 3. As can be seen, the profiles are similar to those of the HN solution presented in Figure 1. The ideal WS solution, however, presents less volatile flow profiles than those of the HN counterpart. This is because inventory can be planned differently for each scenario (we have perfect information). This is particularly evident in the low and high demand scenarios. The total power consumed in the WS scenarios are 88245 kWh, 98914 kWh, and 112779 kWh, respectively. This is less than 1% per scenario, compared with the HN solution. Clearly, significant resiliency can be gained by using the HN formulation over the deterministic one without sacrificing much performance over the ideal case of perfect information. These results illustrate that significant improvements in resiliency can be obtained by anticipating gas demand scenarios of power plants. Unfortunately, *deriving uncertainty characterizations for power plant demands is extremely difficult* because these are fully correlated to the decisions of the ISO on the power grid side. Consequently, coordination between gas and power grid sides would require the ISO to provide suitable demand scenarios that are compatible with its dispatch operations. We now use a case study in Illinois to illustrate this point and other economic issues. We also use this larger system to quantify the effect of increased control flexibility achieved by coordination.

The Illinois power grid transmission system comprises 2,522 lines, 1,908 nodes, 870 demands points, and 225 generators points (153 gas-fired generators). Because




Fig. 4 Illinois electrical (thin lines) and gas (thick lines) transmission systems. Black dots are gas-fired power plants.

of the difficulty in obtaining natural gas infrastructure data, we construct a simulated natural gas network system using the basic topology reported by the EIA⁴ and by using engineering insight to ensure gas supply to all the gas-fired power plants under nominal conditions. The gas network comprises 215 pipeline segments, 157 nodes, 12 compression stations, and 4 supply points. The resulting network is sketched in Figure 4.

Table 1 Economic performance under coordinated and uncoordinated settings (scm= standard cubic meters and M\$=million U.S. dollars).

	ϕ_{grid} [M\$]	ϕ_{gas} [M\$]	$\phi^{gas,comp}$ [\$]	$d^{gas,target}$ [scm $\times 10^{-6}$]	d^{gas} [scm $\times 10^{-6}$]	\mathcal{R} [M\$]
Uncoord	36.54	-13.52	28,618	141.25	135.54	2.70
Coord	36.40	-14.54	33,600	145.74	145.74	3.50

We compare economic performance for the infrastructures under coordinated and uncoordinated settings. The results are summarized in Table 1. In our simulations, the electrical loads were always satisfied; consequently, we report only the generation cost component of the grid cost (we denote this as ϕ_{grid}). From the results we make the following observations:

⁴ <http://tinyurl.com/cssgxr9>

- Under a coordinated setting the power cost decreases by 0.38% which represents a total of \$140,000. The gas cost decreases by 7.54%, which corresponds to a total of \$1,020,000.
- Under an uncoordinated setting only 96% of the gas requested is delivered. At a gas price of 3 \$/MMBTU⁵, the total undelivered gas has an economic value of \$605,000.
- Under a coordinated setting the compression cost increases by 17.4%. This is the result of an increased amount of gas delivered to the power plants. In particular, 7.5% more gas is delivered under the coordinated setting. At a gas price of 3 \$/MMBTU, the value of the additional gas delivered is \$1,080,000. Note that the total increase in compression cost is negligible compared to the additional value of the delivered demand.
- Under a coordinated setting the revenue for the gas-fired generators increases by 29.6%, which corresponds to a total of \$800,000. This is the result of the additional gas delivered and the decreased revenue penalties resulting from coordination.

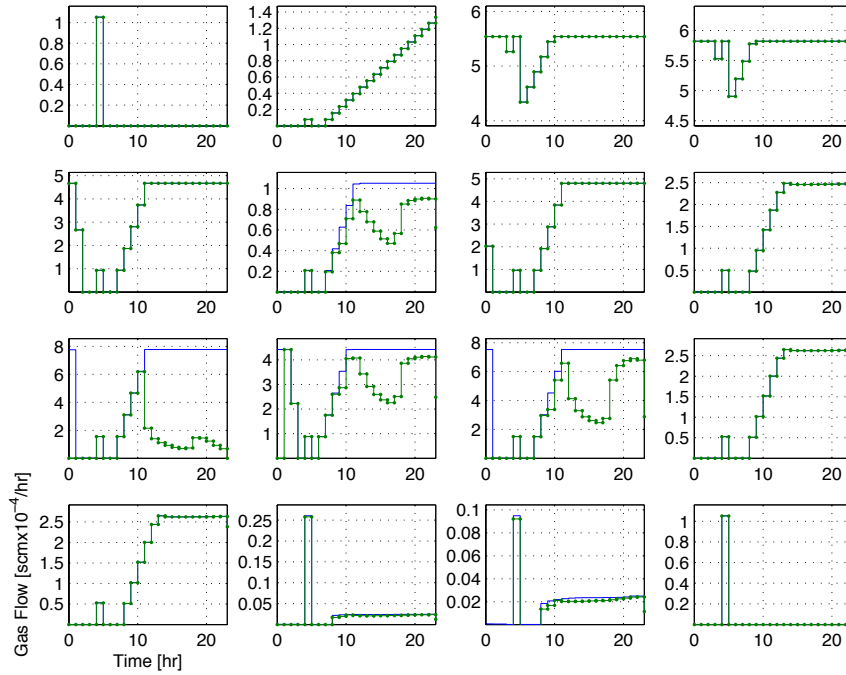


Fig. 5 Requested (blue solid line) and realized (green dotted line) gas demands for 16 power plants under an uncoordinated setting.

⁵ <http://www.eia.gov/naturalgas/weekly/>

It is rather surprising that both the gas and power grid sides benefit from coordination (i.e., *the objectives of the gas and grid operators do not compete*). Moreover, gas-fired generators increase their revenue. We can explain the decreased performance under an uncoordinated setting from the fact that the power grid operator cannot easily determine how much gas can the gas network deliver at different spatial locations and at different times. Thus, the power grid operator can be overly optimistic (as in the case presented) or pessimistic about the amount of gas that can actually be delivered. This situation is clearly illustrated in Figure 5, where we present the target and realized gas demands for 16 different gas-fired generators under the uncoordinated setting. Note that the gas network cannot deliver the total amount of gas requested at four locations. The resulting error in the prediction introduces a penalty for both the power grid and the gas-fired generators. In particular, the power grid operator has to dispatch more expensive power plants, resulting in higher a generation cost and the gas-fired power plants have to pay for the unserved generation. Note also that, *even if the gas operator knows the gas demands of the power grid in advance*, it cannot guarantee to satisfy such demands due to physical constraints.

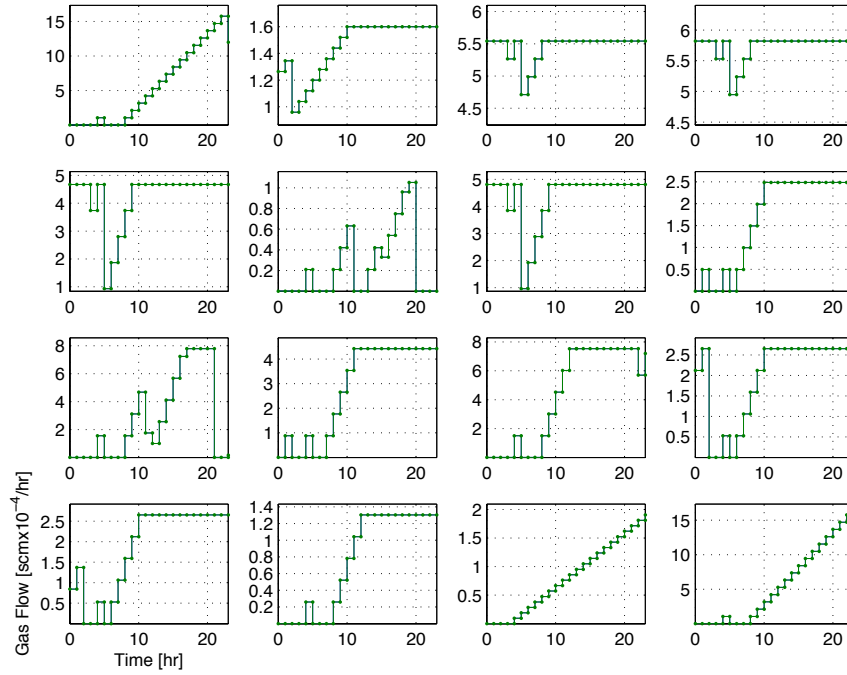


Fig. 6 Requested (blue solid line) and realized (green dotted line) gas demands for 16 power plants under coordinated setting.

The presence of gas shortages at several power plant locations *would suggest that the gas pipeline system is constrained by design*. We now demonstrate that this perception is not correct, under a coordinated setting this *delivery bottleneck can in fact be avoided through better control*. In fact, as we have already seen in Table 1, *7% more gas can be delivered under the coordinated setting*; we now explain why this is the case. From Figure 6 we see that dispatch under a coordinated setting is smoother (i.e., demands are ramps instead of aggressive withdrawals) and this significantly enhances the flexibility to deliver gas at other locations. In particular, all gas demands can be delivered. We can explain this increased flexibility by noticing that, under a coordinated setting, gas-fired generators act as *distributed demand response resources* that the gas operator can use to better control network pressures and flows and avoid delivery bottlenecks. In other words, *gas-fired power plants become assets rather than liabilities to the pipeline operator*.

Figures 5 and 6 also help us illustrate the rather arbitrary patterns of gas demands emanating from power plants, which are dictated by dispatch decisions of the ISO. These gas demands patterns are to be contrasted with gas and electricity demands of urban areas, which follow strong periodic patterns that are correlated to weather and behavior. The irregular patterns of power plants demands make it difficult to model their uncertainty and this can prevent the adoption of advanced stochastic optimal control solutions. The ISO plays a critical role here, as it can provide characterizations of uncertainty by using possible gas demand scenarios generated under different wind power, weather, and contingency scenarios.

4 Computational Issues

In this section we review problem structures arising in emerging optimal control models and discuss available solution strategies. We also highlight limitations of state-of-the-art solvers and discuss open challenges.

4.1 Emerging Model Structures

The deterministic (single scenario) grid-gas OCP can be expressed in the more general form

$$\min \sum_{p \in \mathcal{P}} \varphi_p(x_p) \quad (22a)$$

$$\text{s.t. } c_p(x_p) = 0, p \in \mathcal{P} \quad (\lambda_p) \quad (22b)$$

$$x_p \geq 0, p \in \mathcal{P} \quad (v_p) \quad (22c)$$

$$\sum_{p \in \mathcal{P}} \Pi_p x_p = 0 \quad (\lambda_0) \quad (22d)$$

where $\mathcal{P} := \{gas, grid\}$. Here, $\lambda_0 \in \mathfrak{R}^{m_0}$ are the multipliers of the coupling constraints, $\lambda_p \in \mathfrak{R}^{m_p}$, $v_p \in \mathfrak{R}^{n_p}$ are the multipliers of the partition p constraints and bounds, $x_p \in \mathfrak{R}^{n_p}$ are the primal variables of the partition p , and $\Pi_p \in \mathfrak{R}^{m_0 \times n_p}$ are the coupling matrices. It is important to observe that this structure also arises in stochastic programming problems. For instance, the stochastic OCP for the gas side or for the power grid side can be expressed as (22).

Under a primal-dual interior point framework, the KKT system of problem (22) can be permuted into the block-bordered-diagonal (BBD) form:

$$\underbrace{\begin{bmatrix} K_0 & B_1^T & B_2^T & \dots & B_{|\mathcal{P}|}^T \\ B_1 & K_1 & & & \\ B_2 & & K_2 & & \\ \vdots & & & \ddots & \\ B_{|\mathcal{P}|} & & & & K_{|\mathcal{P}|} \end{bmatrix}}_{:=M(\delta_w, \delta_c)} \begin{bmatrix} \Delta w_0 \\ \Delta w_1 \\ \Delta w_2 \\ \vdots \\ \Delta w_{|\mathcal{P}|} \end{bmatrix} = - \begin{bmatrix} r_0 \\ r_1 \\ r_2 \\ \vdots \\ r_{|\mathcal{P}|} \end{bmatrix}, \quad (23)$$

where $\Delta w_0 = \Delta \lambda_0$, $\Delta w_p = (\Delta x_p, \Delta \lambda_p)$,

$$K_0 = -\delta_c \mathbb{I}_{m_0}, \quad K_p = \begin{bmatrix} W_p(\delta_w) & J_p^T \\ J_p & -\delta_c \mathbb{I} \end{bmatrix} \\ B_p^T = [\Pi_p \ 0], \quad (24)$$

$J_p = \nabla_{x_p} c_p(x_p)$, $W_p(\delta_w) = \nabla_{x_p, x_p} \mathcal{L} + X_p^{-1} V_p + \delta_w \mathbb{I}$. Symbol \mathcal{L} denotes the Lagrange function of (22) and $\delta_w, \delta_c \geq 0$ are regularization parameters. For more details in the derivation of the KKT system the reader is referred to [15]. The BBD structure of the linear system can be exploited using a parallel Schur decomposition of the form

$$\underbrace{-\left(\delta_c \mathbb{I}_{m_0} + \sum_{p \in \mathcal{P}} B_p K_p^{-1} B_p^T\right)}_S \Delta w_0 = -r_0 + \sum_{p \in \mathcal{P}} K_p^{-1} B_p r_p \quad (25a)$$

$$K_p \Delta w_p = -r_p - B_p^T \Delta w_0, \quad p \in \mathcal{P}, \quad (25b)$$

where S is the Schur complement. One can exploit individual structures within each partition. For instance, we can express the optimization problem in the gas partition $p = gas$ as,

$$\min \varphi(x, u) \quad (26a)$$

$$\text{s.t. } c(x, u) = 0, \quad (\lambda) \quad (26b)$$

$$x \geq 0 \quad (v_x) \quad (26c)$$

$$u \geq 0, \quad (v_u). \quad (26d)$$

In this structure, u is assumed have the same dimension as the number of degrees of freedom of the problem (e.g., boost pressures and demands). In other words, if u is fixed, then $c(x, u) = 0$ is a square system of equations (in the gas system these correspond to the discretized transport and network equations). The KKT system of problem (26) is given by

$$\begin{bmatrix} W_{xx}(\delta_w) & W_{xu} & J_x^T \\ W_{ux} & W_{uu}(\delta_w) & J_u^T \\ J_x & J_u & \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta u \\ \Delta \lambda \end{bmatrix} = - \begin{bmatrix} r_x \\ r_u \\ r_\lambda \end{bmatrix} \quad (27)$$

where the coefficient matrix on the left hand side is also known as the augmented matrix. By construction, the Jacobian $J_x = \nabla_x c(x, u)$ is square; and, if it is nonsingular and δ_c , then we can construct the following null-space matrix:

$$Z = \begin{bmatrix} -J_x^{-1} J_u \\ I \end{bmatrix}. \quad (28)$$

The step for u can then be obtained by solving a reduced system of the form $Z^T W(\delta) Z \Delta u = r_Z$ where r_Z is an appropriate right-hand side vector and $Z^T W(\delta) Z$ is the reduced Hessian. Having the step for u , we compute the step for x from $\Delta x = -J_x^{-1}(r_u + J_u \Delta u)$. Note that this approach requires factorizations of J_x and of the reduced Hessian $Z^T W(\delta) Z$ instead of factorizations of the entire augmented matrix. Because of this, this approach can yield significant speed-ups when the number of degrees of freedom u is small.

The expression of structures also facilitates model construction and can accelerate model processing overhead of algebraic modeling languages (e.g., generation of derivative information and sparsity structures) which is significant in large-scale models such as the ones arising in large infrastructures. For instance, the gas-electric coupled problem when discretized in space gives an NLP with 249,919 variables, 224,292 equality constraints, and 154,093 inequality constraints. If we replicate this model over multiple scenarios we can see that 100 scenarios already give an NLP with 25 million variables that cannot be processed with existing algebraic modeling languages such as AMPL, GAMS, or JuMP. Consequently, it is necessary to partition the problem to enable processing and model storage in memory.

We now demonstrate the benefits of identifying structures and exploiting them in high-performance computers. To do so, we use the parallel interior point solver PIPS-NLP [11] to solve a stochastic OCP for a gas pipeline system with 13 nodes and 12 pipelines and 10 compressors. We consider two settings: in the first case we exploit only the stochastic structure (Table 2) while in the second case we exploit the stochastic and the reduced space structure of the problem (Table 3). The NLP under study has 96 scenarios and a total of 1,930,752 variables. The problem is solved on the distributed-memory cluster Fusion at Argonne National Laboratory. As can be seen, strong scaling is observed in both cases and solution times can be brought down from an hour to less than 7 minutes. Moreover, these NLPs cannot be solved in a single processor because of the large memory requirements. We also

observe that exploiting the reduced space structure inside each scenario decreases the computational time by a factor of nearly 3.

Table 2 Scalability of PIPS-NLP exploiting stochastic structure.

Scenarios	n	Obj	Iter	Time(hh:mm:ss)	MPI Proc
96	1,930,752	1.39×10^2	42	01:13:16	8
96	1,930,752	1.39×10^2	42	00:38:18	16
96	1,930,752	1.39×10^2	42	00:24:55	24
96	1,930,752	1.39×10^2	42	00:19:23	32
96	1,930,752	1.39×10^2	42	00:12:42	48
96	1,930,752	1.39×10^2	42	00:06:48	96

Table 3 Scalability of PIPS-NLP exploiting stochastic and reduced space structure.

Scenarios	n	Obj	Iter	Time(hh:mm:ss)	MPI Proc
96	1,930,752	1.39×10^2	42	00:29:54	8
96	1,930,752	1.39×10^2	42	00:14:45	16
96	1,930,752	1.39×10^2	42	00:10:00	24
96	1,930,752	1.39×10^2	42	00:07:36	32
96	1,930,752	1.39×10^2	42	00:05:14	48
96	1,930,752	1.39×10^2	42	00:02:54	96

4.2 Dealing with Negative Curvature

A key difference between convex and nonconvex NLPs is the potential presence of *negative curvature*. The presence of negative curvature indicates that the Newton step computed from the solution of the KKT system (23) might not correspond to a minimum of the associated quadratic programming problem. In a line-search setting this is an important issue because the Newton step cannot be guaranteed to provide a descent direction for the objective function when the constraint violation is sufficiently small, which is a key requirement to ensure global convergence.

The presence of negative curvature can be checked by computing the inertia of the KKT system (23). This can be done using Haynsworth's formula:

$$\text{Inertia}(M(\delta_w, \delta_c)) = \sum_{p \in \mathcal{P}} \text{Inertia}(K_p) + \text{Inertia}(S). \quad (29)$$

We recall that $n = \sum_p n_p$ and $m = m_0 + \sum_p m_p$. Consequently, if we have that $\text{Inertia}(K_p) = \{n_p, m_p, 0\}$ for all $p \in \mathcal{P}$ then the inertia of $M(\delta_w, \delta_c)$ is correct if and only if $\text{Inertia}(S) = \{0, m_0, 0\}$. When the inertia is correct we can thus guarantee that

the Newton step is a descent direction. One can obtain the inertia of the blocks K_p using LBL^T factorizations. If the problem has nested structures, one can also obtain the inertia of each block K_p by applying Haynsworth's formula recursively.

If the inertia of the linear system (23) (or a subsystem) obtained from (29) is not correct we progressively increase the regularization parameter δ_w until the KKT system has correct inertia. We call this the *inertia-based* regularization strategy. Note that every time we increase the regularization parameter we need to solve the linear system again. Consequently, the presence of negative curvature can make the computation of the Newton step very expensive.

Obtaining the inertia of Schur matrices using an LBL^T factorization can be inefficient because the Schur complement is often a dense matrix or contains dense blocks. To avoid these limitations, we have recently proposed an inertia-free test of the form,

$$\Delta w^T W(\delta_w) \Delta w \geq \kappa \Delta w^T \Delta w \quad (30)$$

for $\kappa > 0$ [12]. Here, Δw is a Newton step computed along the null-space of the constraint Jacobian and $W(\delta_w)$ is the *entire* Hessian matrix of the NLP (22). In the inertia-free approach we increase the regularization parameter until the curvature test (30) holds. The key observation is that this test implicitly requires the Newton step to be a descent direction directly which can occur even if the linear system does not have correct inertia. Consequently, the inertia-free test provides more flexibility to accept steps. To demonstrate this feature, we solve a large-scale stochastic optimal control instance and compare inertia-based regularization (IBR) and inertia-free regularization (IFR). This is an NLP with 128 uncertain scenarios, 1,024,651 variables, and 1,023,104 constraints. The results obtained with PIPS-NLP are presented in Table 4. We can see that, despite the high nonlinearity of the large-scale instances, the inertia-free approach IFR converges in all instances. In general IFR requires more iterations than does IBR but the number of factorizations is reduced, resulting in faster solutions. We can thus see the effect of adding flexibility to step acceptance is beneficial.

Table 4 Performance of inertia-based and inertia-free regularization strategies on stochastic OCP for gas system.

#MPI	Obj	IBR			IFR		
		Iter	Linear Solves	Time(s)	Iter	Linear Solves	Time(s)
8	1.26E-02	153	278	832	93	106	491
16	1.26E-02	136	251	363	109	122	315
32	1.26E-02	146	274	209	99	112	143
64	1.26E-02	157	286	123	101	114	79
128	1.26E-02	145	275	64	109	125	52

4.3 Open Issues

We can pose the stochastic programming problem for the coupled grid-gas system (21a) in the general form:

$$\min \sum_{p \in \mathcal{P}} \sum_{j \in \mathcal{P}_p} \varphi_{p,j}(x_{p,j}) \quad (31a)$$

$$\text{s.t. } c_{p,j}(x_{p,j}) = 0, \quad p \in \mathcal{P}, \quad j \in \mathcal{P}_p \quad (\lambda_{p,j}) \quad (31b)$$

$$x_{p,j} \geq 0, \quad p \in \mathcal{P}, \quad j \in \mathcal{P}_p \quad (v_{p,j}) \quad (31c)$$

$$\sum_{j \in \mathcal{P}_p} \Pi_{p,j} x_{p,j} = 0 \quad (\lambda_{0,j}) \quad (31d)$$

$$\sum_{p \in \mathcal{P}} \Pi_p x_p = 0 \quad (\lambda_0) \quad (31e)$$

Here, \mathcal{P} is the scenario set and $\mathcal{P}_p = \text{gas}, \text{grid}$ is the system partition. One can show that this nested problem yields an augmented system of the form (23) in which each diagonal block K_p has a BBD structure of the form

$$K_p = \begin{bmatrix} K_{p,0} & B_{p,1}^T & B_{p,2}^T & \cdots & B_{p,P_p}^T \\ B_{p,1} & K_{p,1} & & & \\ B_{p,2} & & K_{p,2} & & \\ \vdots & & & \ddots & \\ B_{p,P_p} & & & & K_{p,P_p} \end{bmatrix}, \quad p \in \mathcal{P}. \quad (32)$$

Consequently, we can also apply a Schur decomposition to perform solves with the block system K_p . Moreover, structures such as the reduced-space structure of the gas side can be exploited as well. This would give a linear system with 3 nested structures.

The coupled stochastic problem is of interest because it would be desirable to understand the increasing resiliency gained by coordination. As we have seen, for instance, coordination enables the gas infrastructure to deliver significantly larger amounts of gas and this flexibility can be used to withstand abrupt variations of wind power. Solving coupled stochastic problems on realistic networks, however, is extremely challenging and defies the scope of state-of-the-art solvers. To give an idea of the complexity, for the Illinois system we have found that the solution time of a single scenario problem for the coupled gas-electric system is 40 minutes. While it is possible to partition the Illinois grid and gas systems using Schur decomposition, this is not always beneficial. In particular, we have found that performing a direct sparse factorization of the entire coupled system is more efficient than performing Schur decomposition. In other words, the benefits of Schur decomposition are only observed when the spatial domain (network size or spatial discretization resolution) of the gas and network systems increase. Consequently, the only alternative to accelerate solutions seems to coarsen the spatial discretization of the gas transport equations.

To illustrate the effect *coarsening*, we compare the economic and computational performance of coupled gas-electric problems with low- and high-resolution spatial discretizations. We compare the results using our base implementation with $N_x = 10$ spatial points per pipeline and a low resolution implementation with $N_x = 3$ spatial points per pipeline. The low-resolution problem gives an NLP with 141,559 variables, 115,932 equality constraints, 157,765 inequality constraints, and 25,627 degrees of freedom. The high-resolution problem gives an NLP with 249,919 variables, 224,292 equality constraints, 154,093 inequality constraints, and 25,627 degrees of freedom. The number of degrees of freedom remains unchanged because all of these enter at the network nodes and are thus independent of the discretization resolution [27]. This is an important structural property of gas optimal control formulations. The results comparing high and low resolutions are presented in Table 5. The solution time is reduced from 40 minutes to about 10 minutes. Most notably, from Table 6 we can see that coarsening does not introduce large errors in economic performance. This behavior, however, cannot be guaranteed in general. It is thus necessary to devise linear algebra strategies that can perform coarsening adaptively at the linear algebra level to create preconditioners. Multigrid schemes are available for solving OCPs with embedded PDEs but these schemes are currently not general enough to handle hyperbolic PDEs as those arising in gas networks and to handle mixed sets of constraints. This is an important research area.

Table 5 Computational results for coupled problems for base and perturbed topologies.

	Iter.	Time(s)	Linear Solve [-]	Time/Linear Solve(sec)
$N_x = 10$	232	2401.01	311	7.72
$N_x = 3$	157	573.68	188	3.05

Table 6 Economic performance under low- and high-resolution spatial discretizations.

	φ_{grid} [M\$]	φ_{gas} [M\$]	$\varphi_{gas,comp}$ [\$]	$d_{gas,target}^{gas}$ [scm $\times 10^{-6}$]	d_{gas} [scm $\times 10^{-6}$]	\mathcal{R}_{gas} [M\$]
$N_x = 10$	36.40	-14.54	33,600	145.74	145.74	3.50
$N_x = 3$	36.39	-14.55	33,356	145.83	145.83	3.48

5 Conclusions

We have discussed characteristics of emerging optimal control models for interconnected natural gas and electrical networks. These models are motivated by the increasing interest in understanding the economic benefits of coordination and the need to understand system resiliency in the face of extreme weather events and renewable power adoption. We have seen that optimal control problems arising in

infrastructures are highly structured and that these structures can be exploited to accelerate solutions and avoid memory bottlenecks. We have also demonstrated that existing state-of-the-art tools can handle large and highly nonlinear models but these capabilities are insufficient to handle geographical regions of practical interest.

Acknowledgements This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357. Victor M. Zavala acknowledges the support of the Early Career program of the Department of Energy. We also acknowledge the computing resources provided on Fusion and Blues, high-performance computing clusters operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

References

1. Report from growing concerns, possible solutions: The interdependency of natural gas and electricity systems. MITEI Technical Report (2013)
2. January 2014 FERC data request. ISO New England Technical Report (2014)
3. Abbaspour, M., Chapman, K.: Nonisothermal transient flow in natural gas pipeline. *Journal of Applied Mechanics* **75**(3), 031,018 (2008)
4. An, S., Li, Q., Gedra, T.W.: Natural gas and electricity optimal power flow. In: *Transmission and Distribution Conference and Exposition, 2003 IEEE PES*, vol. 1, pp. 138–143. IEEE (2003)
5. Arnold, M., Negenborn, R.R., Andersson, G., De Schutter, B.: Distributed predictive control for energy hub coordination in coupled electricity and gas networks. In: *Intelligent Infrastructures*, pp. 235–273. Springer (2010)
6. Baumrucker, B., Biegler, L.T.: MPEC strategies for cost optimization of pipeline operations. *Computers & Chemical Engineering* **34**(6), 900–913 (2010)
7. Biegler, L.T.: *Nonlinear programming: concepts, algorithms, and applications to chemical processes*. Society for Industrial and Applied Mathematics (2010)
8. Carter, R.G., Dupont, T.F., Rachford Jr, H.H.: Pack management and transient optimization of natural gas transmission lines. In: *IGRC Conference, Vancouver* (2004)
9. Centolella, P., Ott, A.: The integration of price responsive demand into PJM wholesale power markets and system operations. HEPG Technical Report (2009)
10. Chaudry, M., Jenkins, N., Strbac, G.: Multi-time period combined gas and electricity network optimisation. *Electric Power Systems Research* **78**(7), 1265 – 1279 (2008)
11. Chiang, N., Petra, C.G., Zavala, V.M.: Structured nonconvex optimization of large-scale energy systems using PIPS-NLP. In: *Proc. of the 18th Power Systems Computation Conference (PSCC)*, Wroclaw, Poland (2014)
12. Chiang, N., Zavala, V.M.: An inertia-free filter line-search algorithm for large-scale nonlinear programming (2014). Preprint ANL/MCS-P5197-0914, Argonne National Laboratory
13. Chiang, N., Zavala, V.M.: Large-scale optimal control of interconnected natural gas and electrical transmission systems (2015). Preprint ANL/MCS-P5348-0515, Argonne National Laboratory
14. Geidl, M., Andersson, G.: Optimal power flow of multiple energy carriers. *Power Systems, IEEE Transactions on* **22**(1), 145–155 (2007)
15. Kang, J., Chiang, N., Laird, C.D., Zavala, V.M.: Nonlinear programming strategies on high-performance computers. In: *Proc. of the IEEE Conference on Decision and Control, Osaka, Japan* (2015)
16. Liu, C., Shahidehpour, M., Fu, Y., Li, Z.: Security-constrained unit commitment with natural gas transmission constraints. *Power Systems, IEEE Transactions on* **24**(3), 1523–1536 (2009)

17. Liu, C., Shahidehpour, M., Wang, J.: Coordinated scheduling of electricity and natural gas infrastructures with a transient model for natural gas flow. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **21**(2), 025,102 (2011)
18. Marqués, D., Morari, M.: On-line optimization of gas pipeline networks. *Automatica* **24**(4), 455–469 (1988)
19. Ott, A.L.: Experience with PJM market operation, system design, and implementation. *Power Systems, IEEE Transactions on* **18**(2), 528–534 (2003)
20. Pritchard, G., Zakeri, G., Philpott, A.: A single-settlement, energy-only electric power market for unpredictable and intermittent participants. *Operations Research* **58**(4-part-2), 1210–1219 (2010)
21. Qadrdan, M., Chaudry, M., Wu, J., Jenkins, N., Ekanayake, J.: Impact of a large penetration of wind generation on the GB gas network. *Energy Policy* **38**(10), 5684 – 5695 (2010)
22. Rachford Jr, H.H., Carter, R.G., Dupont, T.F., et al.: Using optimization in transient gas transmission. In: PSIG Annual Meeting. Pipeline Simulation Interest Group (2009)
23. Rachford Jr, H.R., Carter, R.G.: Method and apparatus for determining optimal control settings of a pipeline (2004). US Patent 6,701,223
24. Rawlings, J.B., Angeli, D., Bates, C.N.: Fundamentals of economic model predictive control. In: *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 3851–3861. IEEE (2012)
25. Ríos-Mercado, R.Z., Borraz-Sánchez, C.: Optimization problems in natural gas transportation systems: A state-of-the-art review. *Applied Energy* **147**, 536 – 555 (2015)
26. Steinbach, M.C.: On PDE solution in transient optimization of gas networks. *Journal of Computational and Applied Mathematics* **203**(2), 345–361 (2007)
27. Zavala, V.M.: Stochastic optimal control model for natural gas networks. *Computers & Chemical Engineering* **64**, 103–113 (2014)
28. Zavala, V.M., Biegler, L.T.: Optimization-based strategies for the operation of low-density polyethylene tubular reactors: nonlinear model predictive control. *Computers & Chemical Engineering* **33**(10), 1735–1746 (2009)

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.